AMENDMENTS TO THE SPECIFICATION

On page 1, please replace paragraph [0002] as follows:

[0002] Scene change detection in digitized video is the process of identifying changes in the video data that are likely to represent semantically meaningful changes in the video content, typically associated with a change in [[the]] shot or scene. Scene change detection is useful particularly for video indexing and archiving, and allows for user to "browse" video. Conventional scene change detection is typically applied during the playback of video. Numerous algorithms have been created to identify scene changes, relying variously on histogram analysis of color, luminance, and other image characteristics, and in some cases on analysis of motion vectors. These approaches have been likewise extended to identify scene changes in compressed video content, such as MPEG-2 video. However, these approaches require the computationally expensive process of first decoding the compressed video prior to do the scene change analysis, in order to obtain the necessary statistics and data, such as motion vectors, color data, and the like.

On pages 2-3, please replace paragraph [0005] as follows:

[0005] The present invention includes a scene change identification process that operates during the encoding of compressed video, and that provides for flexible control of the bit rates used during the encoding process. Generally, an uncompressed image is received in an encoder or encoding process. The uncompressed image comprises [[a]] macroblocks, and will have a frame type, such as an I frame, a P frame, or a B frame, as is known in the art. Each macroblock will

also have a type, corresponding to one of these types, resulting from a motion compensation process. From the distribution of the types of the macroblocks, a determination is made as to whether the image corresponds to a scene change. In one embodiment, the determination of whether the image corresponds to a scene change is made based on various percentages of predicted motion macroblocks relative to a total number of macroblocks. The percentage of motion blocks is compared with a threshold to determine whether the frame represents a scene change. Based on this determination and the type of the frame, a scene change is determined. Preferably, when a scene change is detected, the encoder allocates a greater number of bits to the frame, without changing the frame type. The frame is then encoded based on [[the]] rate control parameter. The encoder can also generate a side information file that contains the scene change information to create a scene table of contents.

On page 5, please replace paragraph [0012] as follows:

[0012] Having determined whether the current image is a scene change, the scene change detector 50 provides a scene change flag to the quantizer 70 to indicate a scene change. The quantizer 70 responds to the scene change flag by ~~changing the~~ increasing the number of bits used to quantize the image, thereby increasing the quality of the encoded image. This may be done either directly, by changing the quantization step size, or indirectly by ~~changing~~ adjusting [[the]] a rate controller internal to the quantizer 70. The details of the quantization control are further described below.

On page 6, please replace paragraphs [0013] and [0014] as follows:

[0013] In one embodiment, the scene change detector 50 operates as follows. As noted above, there are three types of frames or pictures that will be received by the encoder: I, P, and B frames. If the current frame is an I frame, and a scene change (as would be perceived by a viewer) has in fact occurred here, there is no need for the scene change detector 50 to instruct the quantizer 70 to allocate additional bits during quantization. This is because the I frame will already be entirely intra-coded, and will provide sufficient image quality. In one embodiment in which the scene change detector 50 builds a scene table of contents (TOC) 130, the scene change detector 50 generates index information that identifies the current image as a scene change, for example by outputting the current group of pictures (GOP) index, and the frame number of the current frame within the current GOP. The scene TOC 103 can be transmitted to, subsequently decoded by a complementary decoder to provide the viewer with access to a scene change table of contents or the like.

[0014] The second case is where the current frame is a B ~~frames~~ frame, which will have two reference frames, one forward and one behind. If a scene change occurs at a B frame, then the motion vectors for that frame should mostly point in one direction, either forward, or backward, and hence the distribution of P and B blocks will indicate a scene change.

On pages 7-8, please replace paragraphs [0023] and [0024] as follows:

[0023] The third case is where the current frame is a P frame. As described above, a P frame is encoded based on a prior I frame. The blocks in the P frame then will be distributed ~~amount~~ among I, S, and P blocks. If a scene change has not occurred at the current P frame, then there should be a high proportion of P blocks, since the will be well predicted by the prior I frame.

On the other hand, a scene change at the P frame would correspond to a significant amount of new material in the frame, as compared to a prior reference frame. This means that for such a scene change, the motion estimator 20 should have <u>been</u> unable to find a significant number of matches for the P frame from the blocks of prior reference frame. As a result, in a P frame, a scene change correlates to a relatively high number of I blocks in the frame (since relatively few blocks were motion predicted).

[0024] More specifically, the scene change detector 50 can implement the foregoing logic in a variety of ways. In one embodiment, the scene change detector 50 determines the percentage of I blocks (PI) relative to M. If PI is greater than 65%, then a scene change is declared. Alternatively, the scene change detector 50 can compute the ratio of I blocks to P blocks, and declare a scene change when this ratio is greater than about 2 (that is, more than about two thirds of the frame are I blocks <u>)</u>. Those of skill in the art will appreciate the variety of ways in which statistics of the distribution of blocks types can be computed and selected in view of the teachings of the present invention.

On page 8, please replace paragraph [0026] as follows:

[0026] Optionally, the scene change detector 50 further generates index information that identifies the current P or B frame as a scene change, by outputting the current ~~group of pictures (GOP)~~ <u>GOP</u> index, and the frame number of the current frame within the current GOP, and updating the scene TOC 130.

On pages 8-9, please replace paragraphs [0028] and [0029] as follows:

[0028] As shown by the control signal path from [[the]] buffer 90 to the quantizer 70, the type of the picture and the amount of data in the buffer 90 control the quantization. In a conventional encoder, if the amount of data in the buffer increases (e.g., when the buffer is holding an image that include significant amounts of intra encoded data), the bit rate control algorithm would increase the quantization step size, resulting in a coarser quantization, which would to decrease the image quality of the current image, even if the current image corresponds to a scene change. The encoder 101 of the present invention however does not suffer from this defect, since the scene change detector 50 can influence the degree of quantization based on the presence or absence of a scene change.

[0029] The quantizer 70 operates as follows, in one embodiment. The quantizer 70 maintains two bit rate variables: T is a variable that tracks the total number of bits allocated for current frame in the GOP. R is a variable that tracks the number of bits remaining after encoding each frame in the GOP. A rate control algorithm 75 in the quantizer 70 uses these variables allocate bits to each frame to be encoded. The rate controller 75 will allow the number of bits allocated to a frame to vary by about +/- 2-3% per frame. The ability of the rate control algorithm 75 to stick to [[the]] <u>a</u> target <u>allocation</u> is a measure of the efficiency of the rate control and encoder.

On pages 10-11, please replace paragraphs [0035], [0036] and [0037] as follows:

[0035] More specifically, increasing the value of R by the encoding factor X results in an increase <u>of</u> the target bit rate allocation Tp or Tb depending on whether the current frame is a P or B frame. For a P frame , T = Tp and for a B frame T = Tb.

Tp and Tb are generally calculated as follows:

$$T_p = \max\left\{\frac{R}{\left(N_p + \frac{N_b K_p X_b}{K_b X_p}\right)}, \frac{bit\_rate}{8 \times picture\_rate}\right\}$$

$$T_b = \max\left\{\frac{R}{\left(N_b + \frac{N_p K_b X_p}{K_p X_b}\right)}, \frac{bit\_rate}{8 \times picture\_rate}\right\}$$

[0036] Where Kp and Kb are "universal" constants dependent on [[the]] quantization matrices, with typical values of Kp = 1.0 and Kb = 1.4; N is the number of pictures in the current GOP; Np and Nb are the number of P-pictures and B-pictures remaining in the current GOP in the encoding order; and bit_rate and picture_rate are previously determined factors.

[0037] From the revised target allocations Tb or Tp, a fullness measure ~~dj^p or dj^b~~ $\underline{d_j^p \text{ or } d_j^b}$ is computed, where:

$$d_j^p = d_0^p + B_{j-1} - \left(\frac{T_p \times (j-1)}{MB\_cnt}\right)$$

$$d_j^b = d_0^b + B_{j-1} - \left(\frac{T_b \times (j-1)}{MB\_cnt}\right)$$

On pages 11-12, please replace paragraphs [0041], [0042] and [0043] as follows:

[0041] and dj is the fullness of the appropriate virtual buffer. Finally, the quantization parameter Qj is used to ~~quantized~~ <u>quantize</u> the DCT coefficients for the current macroblock.

Further details of the quantization process may be found in the documentation for Test Model 5 for the MPEG Software Simulation Group, MPEG Doc. No. MPEG 93/457 (April, 1993).

[0042] In summary ~~then~~, increasing T and R results in an increase in the number of bits allocated to the current type of frame, which in turn decreases the fullness measure for the current frame type as well. This decrease in the fullness measure in turn reduces Q. In an alternative embodiment, T alone is increased by X, which will also have a similar net effect on reducing Q. Similarly, R alone can be increased. In yet another embodiment, R and T can be increased by ~~differing~~ <u>different</u> amounts. In a further embodiment, Q can be reduced directly based on the above equations.

[0043] After the quantization of the DCT coefficients is performed, ~~then~~ the rate control algorithm <u>75</u> subtracts X from T and from R, essentially restoring them to their prior values. R is then updated to reflect the actual number of bits used to encode the current frame. The temporary values T and R do not effect the operation of the inverse quantizer 100.